

The Power of Local Information in PageRank*

Marco Bressan
Dip. Informatica
Sapienza Università di Roma
Roma, Italy
bressan@di.uniroma1.it

Enoch Peserico
Dip. Ing. Informazione
Università di Padova
Padova, Italy
enoch@dei.unipd.it

Luca Pretto
Dip. Ing. Informazione
Università di Padova
Padova, Italy
pretto@dei.unipd.it

Abstract

How large a fraction of a graph must one explore to rank a small set of nodes according to their PageRank scores? We show that the answer is quite nuanced, and depends crucially on the interplay between the correctness guarantees one requires and the way one can access the graph. On the one hand, assuming the graph can be accessed only via “natural” exploration queries that reveal small pieces of its topology, we prove that deterministic and Las Vegas algorithms must in the worst case perform $n - o(n)$ queries and explore essentially the entire graph, independently of the specific types of query employed. On the other hand we show that, depending on the types of query available, Monte Carlo algorithms can perform asymptotically better: if allowed to both explore the local topology around single nodes and access nodes at random in the graph they need $\Omega(n^{2/3})$ queries in the worst case, otherwise they still need $\Omega(n)$ queries similarly to Las Vegas algorithms. All our bounds generalize and tighten those already known, cover the different types of graph exploration queries appearing in the literature, and immediately apply also to the problem of approximating the PageRank score of single nodes.

Keywords: Local algorithms, graph ranking, PageRank, graph exploration models

*This is the full version of the WWW 2013 poster [13]. Last update: July 8th, 2013.

1 Introduction

Suppose one has to compute the relative ranking induced by PageRank on a small set of nodes of a large, hardly accessible graph. Which fraction of the graph must be explored to assess this relative ranking, and to what extent do different exploration primitives influence the exploration cost? We investigate the interplay between algorithms for this “local PageRank problem” and graph exploration models, focusing on the correctness guarantees of the former and on the locality properties of the latter. For algorithms that guarantee always a correct output (i.e. Las Vegas and deterministic algorithms), we show that no “natural” exploration model exempts from exploring almost the entire input graph in the worst case. In particular, we prove lower bounds on the number of nodes to visit that are tighter than all the existing ones, going as far as $n - o(n)$ for graphs on n nodes; and that hold simultaneously under any exploration model, for any single algorithm execution, and for any value of the PageRank scores and of their relative distance. The cornerstone of such a level of generality is the novel notion of *ranking subgraph*, which characterizes the cost of any algorithm returning always a correct result. While we exploit this fact to compute tight worst-case lower bounds, we also show that, in general, even approximating them within a factor $o(\log(n))$ is an NP-complete problem. On the other hand, for algorithms that can return an incorrect output (i.e. Monte Carlo algorithms), we show that the interplay with the exploration model is crucial: while under any “strictly local” exploration model they still undergo worst-case lower bounds similar to those of Las Vegas algorithms, they turn to be very efficient with just a “minimum amount” of non-locality. More precisely we prove that, if one can only visit neighbours of already visited nodes, then a worst-case $\Omega(n)$ lower bound applies, for any possible value of the PageRank scores and of their relative distance. If instead one is allowed to jump to random nodes in the graph, then the exploration cost can be radically reduced. In this case, using just two spare exploration primitives, we build an algorithm that ranks an arbitrary subset of nodes with arbitrarily small error probability by visiting a portion of the graph inversely proportional to the smallest of their scores – thus ranging between $O(1)$ and $O(n)$. By proving (the first) lower bounds for Monte Carlo algorithms holding under any exploration model, we show that our algorithm is optimal also among those employing more informative primitives, except for very small scores: in this case, an even more efficient (and fully sublinear?) local algorithm for computing the PageRank ranking might exist.

While strengthening the existing bounds [7, 14], our results delineate a clear separation in cost between algorithms that guarantee always a correct solution, and algorithms that accept a small probability of error – but only when these latter can access the graph in a “global” fashion. Incidentally, this justifies why all the existing algorithms for local PageRank computations with worst-case cost guarantees have a positive probability of error and rely on random access to the graph. Although devised for the problem of computing locally the PageRank ranking, our lower and upper bounds apply also to the problem of approximating locally the PageRank scores; the two problems appear then essentially equivalent, answering an open question of [7].

Motivations and Related Work. PageRank [15] is a prominent node centrality measure, proving impressively successful in tasks as diverse as crawl seeding [17], spam detection [20], graph partitioning [3], trendsetter identification [27] and bibliometrics [29]. Indeed, PageRank has been rated among the top 10 algorithms in data mining [28], and its large and growing success has attracted a continuous stream of research. While its algebraic properties have been explored intensely in the past (see [5] and [23] for an overview), in the last years the focus has moved to the algorithmic side, and especially on the issues of its parallel, distributed, incremental, and local computation; we address this latter here.

In this paper, by “locality” we mean the possibility of computing the output of a problem by examining a limited part of the input. When the input is a graph, this typically means that an algorithm should explore a small portion of it, starting from a given set of nodes and

obtaining information on the global structure one piece at a time, each time incurring a cost. Locality, in this sense, has been studied since the eighties, especially with regard to classical problems on graphs (see e.g. [18, 25, 22]). More recently, the difficulty to store, “snapshot”, or even access in their entirety massive graphs (such as the web or social networks) has urged to investigate locality as one of the crucial issues for the wide range of computational problems that naturally arise in large graphs; and this has led to the introduction of models and algorithms for solving (global) graph problems in a local fashion. Notably, [12] introduced the “jump and crawl” graph exploration model, along with local algorithms for finding vertices with extreme topological properties (e.g. with high clustering coefficient), while [9, 10] investigated local algorithms for three graph problems – finding the node of maximum degree, a path between any two nodes, and a minimum dominating set.

The problem of estimating locally the PageRank score of a given node was introduced by Chen et al. in [16], together with heuristic exploration strategies which give good score approximations by visiting only a handful of nodes. From the theoretical point of view, however, under their model any algorithm requires to visit a worst-case $\Omega(n)$ nodes, as shown by our lower bounds. In [19], Fogaras et al. indirectly lower bound the number of nodes one must explore to approximate Personalized PageRank scores [21]. Since (as they note) by a simple reduction this allows to compute the standard PageRank scores, our lower bounds also apply; and, indeed, special cases of their bounds match special cases of ours. Avrachenkov et al. [4] describe several variants of Monte Carlo score approximation algorithms based on random walks (on which leverages the algorithm of Section 5); however, as they require visiting $\Omega(n)$ nodes, they are not apt to local computations. The first lower bounds on the cost of locally approximating the score of a node within a factor $1 + \epsilon$ are given by Bar-Yossef et al. in [7], in a scenario where algorithms are not limited to crawl back from the target node (differently from [16]), but still under a specific exploration model. Their $\Omega(n)$ bound for deterministic algorithms and $\Omega(\sqrt{n})$ bounds for Las Vegas and Monte Carlo algorithms hold for a restricted choice of the damping factor α and for $\epsilon \in O(1)$, and do not allow to choose the target node’s score; in contrast, we give (tighter) bounds holding for any α , for any possible value of ϵ and of the scores (even when non-constant functions of n), and under any graph exploration model. In [2] Andersen et al. give a local algorithm, based on back crawling, to approximate the score contribution of each node in the graph towards a given node v within an additive error ϵ_a . This algorithm can be used to compute an approximation of v ’s score, but in the worst case it requires to visit essentially the entire graph – in line with our $\Omega(n)$ lower bound for “strictly local” algorithms. In [14] Bressan et al. provide, under the same model of [7], the first formalization and analysis of the local PageRank ranking problem, which asks to rank a set of nodes in nondecreasing order of PageRank score, allowing ties for scores within a factor $1 + \epsilon$. Their $\Omega(n)$ bounds for deterministic algorithms and $\Omega(\sqrt{n})$ bounds for Las Vegas and Monte Carlo algorithms match those of [7] for the local score approximation problem, leaving open the question if computing the ranking is as hard as computing the score; and, similarly to [7], they hold for $\epsilon \in O(1)$ and do not allow choosing scores arbitrarily. Finally, Borgs et al. [11], in a model that allows random jumps, provide a Monte Carlo algorithm that returns all the nodes with score at least Δ but no node with score below Δ/c , for any given $c > 3$, by visiting only $\tilde{O}(1/\Delta)$ nodes overall.

Organization of the Paper. After a short review of PageRank (Section 2), Section 3 formally introduces graph exploration models and formally defines the problem at hand. Section 4 introduces ranking subgraphs, uses them to prove lower bounds for deterministic and Las Vegas local ranking algorithms under any exploration model, and analyses the computational complexity of determining their size. Section 5 gives lower bounds for Monte Carlo algorithms under different exploration models, and provides an algorithm performing almost optimally under every model even if exploiting only a minimum degree of “non-locality”.

2 PageRank

Let G be an n -node graph with no dangling nodes (i.e. nodes with no outgoing arcs; see the note below). The PageRank score [24] of a node v is defined as:

$$P(v) = \frac{1 - \alpha}{n} \sum_{\tau=0}^{+\infty} \alpha^\tau \sum_{z \in G} \text{inf}_\tau(z, v) \quad (1)$$

where the *damping factor* α is some constant in $(0, 1)$, and the τ -step influence $\text{inf}_\tau(z, v)$ of z on v is the probability that a “random surfer” starting in z , and at each step following an outgoing arc at random, is in v after exactly τ steps. Using $\text{inf}_\tau(z, v)$ to define the *contribution* $P(z, v)$ of z to v , the score of v can be rewritten as:

$$P(v) = \sum_{z \in G} P(z, v) \quad \text{where} \quad P(z, v) = \frac{1 - \alpha}{n} \sum_{\tau=0}^{+\infty} \alpha^\tau \text{inf}_\tau(z, v) \quad (2)$$

When G contains dangling nodes, the scores given by Equation (1) add up to less than 1, and thus do not form a probability distribution; in this case, PageRank is not defined. This is typically solved by preprocessing the graph, adding to each dangling node n outgoing arcs towards each node of G . One can prove (see [8]) that the resulting PageRank scores are a rescaled version of those given by Equation (1) in the original graph, and thus conserve their ratio and the ranking they induce on the nodes. For this reason, even in the presence of dangling nodes, we will use the scores given by Equation (1), which are equivalent to (but easier to compute than) the scores obtained after the preprocessing. Finally, it is worth remarking that the score of a dangling node multiplies by $1/(1 - \alpha)$ when adding a self-loop to it (see [14]).

3 Graph Exploration, Local Algorithms, and Problem Definition

This section introduces graph exploration models and local algorithms (Subsection 3.1), formalizes the local PageRank ranking problem (Subsection 3.2), concluding with remarks on the information available to algorithms (Subsection 3.3).

3.1 Graph Exploration Models

A *graph exploration model* describes how algorithms can access a graph by defining a set of *queries*, where each query receives an input (possibly empty) and reveals a (typically tiny) portion of the graph in output. In the popular *link server* graph exploration model [7], a query accepts a node in input and returns in output a list of all the nodes pointing to it and a list of all the nodes pointed by it. The *jump and crawl* model [12] allows a *jump* query that returns a random node of the graph, and a *crawl* query equivalent to the link server query; in [11], the *crawl* query returns instead a random child of the input node (and this will be our definition of *crawl*). In the more minimalist *edge probing* model (see e.g. [1]), a query allows to check for an arc’s existence. An algorithm obeying the exploration model is called a *local algorithm*, and its goal is to evaluate one or more nodes in terms of some global graph metric using as few queries as possible – the number of queries is the *cost* incurred by the algorithm. Note that, since the number of queries needed to solve a problem should be a measure of its “exploration complexity”, a query should disclose no information about the global arc structure of the graph (see the note in A.1), if not for a small portion near the queried node. Formally, given a graph G , we require that any graph exploration model satisfy the following property.

Property 1. *The output of a query is a (possibly empty) connected subgraph of G . For global queries, the output must not depend on the arcs of G , but can depend on its set of nodes. For local queries, it can depend on the arcs pointing to and from a node u in input.*

Intuitively, while global queries allow to discover remote nodes of the graph, local queries allow to assess its local structure around a given node. This property captures any “natural” graph exploration model, including all those proposed in literature (and those mentioned above). It will turn out that the model at hand is irrelevant in some cases but crucial in others, and that interesting models are built on three types of query:

- $links(u)$ returns all the parents/children of u and their arcs to/from u .
- $crawl(u)$ returns a child of u uniformly at random, or an empty graph if u has no children.
- $jump()$ returns a node of G uniformly at random.

To capture the hardness of generating unknown, but valid IDs from scratch (think of e.g. web URLs), we do not allow algorithms to query undiscovered nodes, which would require “guessing” their ID; instead, nodes should first be discovered via the queries provided by the model.

3.2 Local Computation of PageRank

Given a graph G , a damping factor α , a set of *target nodes* $v_1, \dots, v_k \in G$, and a separation $\epsilon > 0$, the *local PageRank ranking problem* asks to rank the target nodes in nonincreasing order of their PageRank scores using local algorithms. Any ranking is valid for nodes whose score ratio is below $1 + \epsilon$, in which case we say there is a tie; otherwise we say that the nodes are ϵ -separated, and require that the ranking induced by the scores be provided. The goal of a local (PageRank) ranking algorithm is to compute a correct ranking using as few queries as possible. It is easy to see that a local algorithm approximating the PageRank scores within a factor $\sqrt{1 + \epsilon}$ can immediately be turned into a local ranking algorithm for separations greater or equal to ϵ – therefore, any lower bound on the cost of local ranking for a separation ϵ applies to the cost of local $(\sqrt{1 + \epsilon})$ -approximation of the score, and the converse is true for upper bounds. We denote an instance of the problem as $(G, \{v_1, \dots, v_k\})$, omitting ϵ and α for easiness of notation. We write $u \succ v$ if u ranks higher than v . Special attention will be paid to scenarios where the target nodes are among the top ranking in the graph, which are often the only ones of interest.

3.3 Awareness Degree of the Algorithms

The cost incurred by an algorithm may reduce if the algorithm is aware of some global property of the graph (e.g. its size or its diameter). However, assuming any knowledge appears arbitrary, since in many cases only partial information is available (especially in massive or evolving graphs), and this information also varies from case to case. For these reasons, and consistently with the (sometimes tacit) assumptions of previous work [2, 6], we will consider obliviousness as the reference case – the algorithms do not have any kind of awareness about global properties of the graph, but possess only the information provided by the output of the queries.

4 Deterministic and Las Vegas

This section investigates the interplay between graph exploration models and local ranking algorithms that return always a correct ranking (i.e., deterministic and randomized Las Vegas algorithms). The variety of graph exploration models satisfying Property 1 makes it non-trivial to analyse these local ranking algorithms “in general”. Surprisingly, this can be done by analysing the algorithms under a specific, strictly local exploration model that allows only $links()$ queries; in this context, the two simple notions of *visit subgraph* (Subsection 4.1) and *ranking subgraph* (Subsection 4.2) completely characterize the behaviour of any algorithm that returns always a correct solution, and it will turn out that ranking subgraphs give per-instance lower bounds holding also under any other exploration model. Leveraging on these results,

we build worst-case instances that yield tight lower bounds holding simultaneously under any graph exploration model (Subsection 4.3). Finally, we prove that, in general, computing or even approximating these lower bounds is an NP-complete problem (Subsection 4.4).

4.1 Visit Subgraphs

At any instant, the nodes discovered by a generic algorithm can be divided into a “kernel” set of nodes that have been queried (i.e. given in input to a query), and a “frontier” set of nodes that have not been queried but appeared in some query’s output. Under a model that allows only *links()* queries, the algorithm knows all the arcs between kernel nodes, all the arcs between kernel nodes and frontier nodes, and none of the arcs between frontier nodes. This can be formalized in the notion of *visit subgraph compatible with a given graph* (see Figure 1), where a *visit subgraph* describes the general structure of the information collected via *links()*:

Definition 1. A visit subgraph is a graph $H = (V, A)$ where V is partitioned in a kernel and a frontier such that 1) every frontier node has an arc to or from a kernel node and 2) there are no arcs between two frontier nodes.

Definition 2. A visit subgraph H is compatible with a graph G if 1) H is a subgraph of G , 2) the kernel of H induces the same subgraph in H and in G , and 3) the nodes pointing from or to the kernel of H , and the relative arcs, are the same in H and in G .

Intuitively, a visit subgraph H is compatible with a graph G (we also say that H is on G , or that G is compatible with H) if some sequence of *links()* queries on G yields exactly H . We will denote by $\text{kernel}(H)$ the kernel of H and by $\text{frontier}(H)$ its frontier, and we will call *kernel size* of H the cardinality of $\text{kernel}(H)$.




Figure 1: A visit subgraph H (solid arcs and nodes), with $\text{kernel}(H) = \{u, u_1, v\}$ (in bold) and $\text{frontier}(H) = \{u_2, v_1\}$, compatible with two different graphs G' and G'' (left and right, made of all the nodes and arcs included those dashed). For e.g. $\alpha = 0.3$, this visit subgraph is not a ranking subgraph (see Subsection 4.2) for u and v , since $u \succ v$ in G' but $u \prec v$ in G'' .

4.2 Ranking Subgraphs

Since an algorithm holds only the information given by the visit subgraph retrieved so far (see Subsection 3.3), it cannot decide which of the different compatible input graphs it is exploring. Nonetheless, that information may still be enough to deduce the correct ranking of the target nodes – in this case, we say that the visit subgraph is a *ranking subgraph*, which can be defined as follows:

Definition 3. A visit subgraph H is a ranking subgraph for the nodes v_1, \dots, v_k if $v_1, \dots, v_k \in \text{kernel}(H)$ and there do not exist two graphs compatible with H where v_1, \dots, v_k have different relative rankings (however ties, according to the separation ϵ , are broken).

We can now prove that, under a model that allows only *links()* queries, any algorithm that returns always a correct ranking must visit always a ranking subgraph:

Theorem 1. *Let ALG be a deterministic or randomized Las Vegas local ranking algorithm using only *links()* queries, and let H be the visit subgraph visited by any execution of ALG on a given instance $(G, \{v_1, \dots, v_k\})$. Then H is a ranking subgraph for v_1, \dots, v_k .*

By Theorem 1, in the strictly local model allowing only *links()* queries, the cost incurred by any deterministic or randomized Las Vegas algorithm on the instance $(G, \{v_1, \dots, v_k\})$ is lower bounded by the smallest kernel size of all the ranking subgraphs for v_1, \dots, v_k compatible with G . Crucially, we can prove that this holds also under any other model satisfying Property 1, whatever the combination of local and global queries it allows. Formally:

Theorem 2. *Consider an instance $(G, \{v_1, \dots, v_k\})$, and let r^* be the smallest kernel size of all the ranking subgraphs for v_1, \dots, v_k on G . Then, under any graph exploration model satisfying Property 1, any execution of any deterministic or randomized Las Vegas local ranking algorithm performs at least r^* queries on the instance $(G, \{v_1, \dots, v_k\})$.*

By Theorem 2, ranking subgraphs give per-instance lower bounds, independent of the exploration model, on the cost of solving the local ranking problem. It turns out that ranking subgraphs are characterized by two simple properties, which can be expressed in terms of *kernel score*:

Definition 4. *Consider a visit subgraph H and a node $v \in \text{kernel}(H)$. The kernel contribution $P_H(z, v)$ is the contribution of z to v through paths whose nodes (including z) are all in $\text{kernel}(H)$. The kernel score $P_H(v)$ is:*

$$P_H(v) = \sum_{z \in H} P_H(z, v) \quad (3)$$

Thus, $P_H(\cdot)$ is the “standard” PageRank score on the graph H , except that it disregards any contribution coming from or through $\text{frontier}(H)$. Therefore, given H , $P_H(\cdot)$ lower bounds the PageRank score on any graph compatible with H , where this contribution in general exists; and since it can be arbitrarily small (as frontier nodes can have arbitrarily many children), the bound is the tightest possible. We can now use $P_H(\cdot)$ to characterize ranking subgraphs:

Theorem 3. *A visit subgraph H is a ranking subgraph for two kernel nodes $u \succ v$ if and only if:*

1. $P_H(u) \geq P_H(v)/(1 + \epsilon)$
2. $\forall w \in \text{frontier}(H) : \sum_{(w,z) \in H} P_H(z, u) \geq \sum_{(w,z) \in H} P_H(z, v).$

Theorem 3 implies that an algorithm can use $P_H(\cdot)$ to assess if the current visit subgraph H is a ranking subgraph and, in that case, immediately return the correct ranking. Subsection 4.3 exploits all these results to develop worst-case lower bounds on the cost of local ranking; the proofs will make use of the notion of union between visit subgraphs (formalized in A.5).

4.3 Lower Bounds

Leveraging on the concept of ranking subgraph and related results (Subsection 4.2), we are now able to prove lower bounds for deterministic and randomized Las Vegas algorithms that are independent from the graph exploration model. Formally:

Theorem 4. *Choose an integer $k \geq 2$, a damping factor $\alpha \in (0, 1)$, a score function $\Theta(1/x) \leq p(x) \leq \Theta(1)$, and a separation function $\Theta(1/x) \leq \epsilon(x) \leq \Theta(1/p(x)^{\frac{1}{k-1}})$. There exists a family of graphs $\{G_m\}$ where G_m has size $n = \Theta(m)$ and contains k nodes v_1, \dots, v_k such that:*

1. v_1, \dots, v_k have score $\Omega(p(n))$
2. v_1, \dots, v_k are $\Theta(\epsilon(n))$ -separated, i.e. $P(v_{i+1})/P(v_i) \in \Theta(1 + \epsilon(n))$
3. v_1, \dots, v_k are the top k ranking nodes in G_m
4. under any graph exploration model, any execution of any deterministic or randomized Las Vegas local ranking algorithm needs $n(1 - O(p(n)\epsilon(n)(1 + \epsilon(n))^{k-2}))$ queries to solve the instance $(G_m, \{v_1, \dots, v_k\})$.

Theorem 4 strengthens and extends to every exploration model the $\Omega(n)$ lower bound for deterministic algorithms limited to *links()* queries, and the analogous $\Omega(\sqrt{n})$ lower bound for Las Vegas algorithms (see [14]); and it does so also for the problem of computing within a factor $\Theta(\sqrt{1 + \epsilon(n)})$ the scores of the target nodes (whose solution gives the ranking) [7]. Furthermore, when the absolute score separation $p(n)\epsilon(n)(1 + \epsilon(n))^{k-2}$ between v_{k-1} and v_k is in $o(1)$ (see the proof in A.6), the lower bound goes as far as $n - o(n)$, meaning that the algorithm must explore the entire graph, save possibly a vanishingly small portion.

In the next subsection we will prove that computing or even approximating these general, tight lower bounds on generic instances is an NP-complete problem.

4.4 Ranking Subgraphs and Computational Complexity

If we were able to compute the minimal kernel size of all the ranking subgraphs for any given instance, we could obtain lower bounds for any graph and any set of target nodes – including those deriving from real-world scenarios. Unfortunately, as the present subsection shows, this problem is computationally intractable unless $P = NP$.

Given a graph G and two target nodes u and v (the results can clearly be extended to $k \geq 2$ target nodes), we denote by RANKGRAPH the problem of deciding if G is compatible with a ranking subgraph of given kernel size r for $u \succ v$. Formally:

RANKGRAPH(G, u, v, r)

Input: directed graph G , target nodes $u, v \in G$, integer $r > 0$.

Output: YES if there exists a ranking subgraph of kernel size r for $u \succ v$ compatible with G , NO otherwise.

Theorems 5 and 6 prove that RANKGRAPH is NP-hard even when the complexity deals only with, respectively, Property 1 and Property 2 of Theorem 3 (part of the kernel set being already fixed by the other property); thus, the problem appears to be the “superposition” of two NP-hard problems.

Theorem 5. *RANKGRAPH is NP-hard (by reduction from CLIQUE).*

Theorem 6. *RANKGRAPH is NP-hard (by reduction from DOMINATING SET).*

Furthermore, since DOMINATING SET is NP-hard to approximate within $o(\log(|G|))$ [26], and since the reduction of Theorem 6 preserves (within small constant factors) the size of both the input graph and the dominating set, we have:

Lemma 1. *Approximating RANKGRAPH within $o(\log(|G|))$ is NP-hard.*

Note that verifying if a given visit subgraph is a ranking subgraph for $u \succ v$ in G takes polynomial time – one need only check that its kernel set contains u and v , and that it satisfies both properties of Theorem 3 – and hence RANKGRAPH is NP-complete, even accepting an approximation factor $o(\log(|G|))$.

5 Monte Carlo Algorithms

This section investigates the interplay between graph exploration models and Monte Carlo local ranking algorithms. It turns out that the choice of the exploration model is crucial: depending on the type of queries available, an algorithm can perform as poorly as possible or instead very efficiently. We start by showing that, differently from deterministic and Las Vegas algorithms, Monte Carlo local ranking algorithms cannot be characterized using the theoretical machinery of ranking subgraphs (Subsection 5.1); and nonetheless, under any “strictly local” graph exploration model, they still undergo essentially the same lower bounds (Subsection 5.2). On the other hand, with just a “minimum amount” of non-locality, we are able to build an efficient, random walk-based Monte Carlo local ranking algorithm (Subsection 5.3), which we pair with almost matching lower bounds that hold under all graph exploration models (Subsection 5.4).

5.1 Eluding Ranking Subgraphs

Since ranking subgraphs consist of portions of the graph an algorithm must visit to guarantee always a correct result, one may think that Monte Carlo algorithms can in part avoid them. Indeed, it turns out that Monte Carlo algorithms can elude ranking subgraphs in a strong sense: we show that, even under an exploration model allowing only local queries, there exists one that always solves correctly a particular instance and yet never visits any of its ranking subgraphs, while guaranteeing an arbitrarily small error probability on every other instance. Formally:

Theorem 7. *For any $\eta > 0$ there exist a $\text{links}()$ -based Monte Carlo local ranking algorithm with confidence $1 - \eta$ and an instance $(G, \{u, v\})$ such that the algorithm always solves correctly $(G, \{u, v\})$ without visiting any of its ranking subgraphs.*

At this point, one might wonder if Monte Carlo algorithms can outperform deterministic and Las Vegas algorithms. As long as the exploration model allows only local queries, however, essentially the same bounds hold, and any Monte Carlo algorithm with non-trivial confidence may have to visit almost the entire input graph, as the next subsection shows.

5.2 The Limited Power of Local Information

We prove that, under any graph exploration model that does not allow global queries (thus forcing the exploration to expand from the target nodes via local queries), in the worst case one must visit essentially the entire input graph even to perform just slightly better than returning a ranking at random. Formally:

Theorem 8. *Choose an integer $k \geq 2$, a damping factor $\alpha \in (0, 1)$, a separation $\epsilon > 0$, and a score function $\Theta(1/x) \leq p(x) \leq \Theta(1)$. Under any graph exploration model allowing only local queries, for any Monte Carlo local ranking algorithm MC with confidence $\frac{1}{k!} + \delta$ there exists a family of graphs $\{G_m\}$ where G_m has size $n = \Theta(m)$ and contains k nodes v_1, \dots, v_k such that:*

1. $P(v_1), \dots, P(v_k) \in \Theta(p(n))$
2. v_1, \dots, v_k are $\approx \epsilon$ -separated, i.e. $P(v_{i+1})/P(v_i) \approx 1 + \epsilon$
3. MC performs $\Omega(\delta n)$ queries on the instance $(G, \{v_1, \dots, v_k\})$.

By Theorem 8, as long as we are confined to local queries, the simultaneous presence of randomization and of a positive and possibly large probability of error yields almost no advantage over deterministic algorithms (that cannot perform random choices and can never return an incorrect ranking). As the next subsection shows, this is the borderline: adding just a “minimum amount” of non-locality allows to build much more efficient algorithms, even if the local queries available are far less informative than e.g. $\text{links}()$.

5.3 SampleRank

Under a model allowing only *jump()* and *crawl()* queries, we build an efficient, random walk-based Monte Carlo algorithm for the local computation of PageRank ranking, with guarantees on both the confidence and the cost. The algorithm stems from existing random walk-based techniques for the local approximation of scores, and estimates the ranking behaving optimally (see the bounds of [11]) for every possible value of the nodes' scores and separation (unlike e.g. that of [11], which is designed to work only for $\epsilon > 2$).

Our algorithm is based on the following node sampling routine, that emulates a random walk using *jump()* and *crawl()* queries. As already proven in [4], the routine returns node v with probability exactly $P(v)$; Lemma 2 gives a proof which is more insightful in this context.

Algorithm 1 SampleNode()

```

current_node  $\leftarrow$  jump()
loop
  with probability  $(1 - \alpha)$  return current_node
  current_node  $\leftarrow$  crawl(current_node)
  if current_node =  $\emptyset$  then current_node  $\leftarrow$  jump()
end loop

```

Lemma 2. *SampleNode() returns node v with probability equal to $P(v)$.*

By a simple Chernoff bound, m calls to SampleNode perform approximately $m/(1 - \alpha)$ queries (this could be reduced to m , but making the algorithm more complicated to analyse [4]):

Lemma 3. *One call to SampleNode performs in expectation less than $\frac{2}{1-\alpha}$ queries. The probability that m calls to SampleNode perform more than $\frac{2(1+\Delta)m}{1-\alpha}$ queries is less than $e^{-\frac{m}{2} \cdot \frac{\Delta^2}{1+\Delta}}$.*

By Lemma 2, we can estimate the PageRank score of a node v as the fraction $\hat{P}_m(v)$ of m calls to SampleNode that returned v , and we can in turn estimate the relative ranking of a set of nodes by repeatedly calling SampleNode until their estimated scores “stabilize”. This intuition is formalized in our algorithm SampleRank, which repeatedly samples the nodes of G until, according to confidence intervals based on \hat{P}_m , each pair of target nodes either appears to be ϵ -separated or appears to yield a tie.

Algorithm 2 SampleRank($G, v_1, \dots, v_k, 1 - \eta$)

```

repeat
  perform the  $m$ -th call to SampleNode
   $\hat{P}_m(v_i) \leftarrow$  fraction of SampleNode calls that returned  $v_i$ 
   $C_m(v_i) = (P_m^L(v_i), P_m^U(v_i)) \leftarrow (1 - \eta/k)$ -confidence interval for  $P(v_i)$  based on  $\hat{P}_m(v_i)$ 
until  $\forall i \neq j : (C_m(v_i) \cap C_m(v_j) = \emptyset) \vee (P_m^U(v_i)/P_m^L(v_j) \leq 1 + \epsilon) \vee (P_m^U(v_j)/P_m^L(v_i) \leq 1 + \epsilon)$ 
return the ranking of  $v_1, \dots, v_k$  induced by  $\hat{P}_m$ 

```

We prove that:

Theorem 9. *Consider an instance $(G, \{v_1, \dots, v_k\})$, and let $p = \min_{i=1, \dots, k} \{P(v_i)\}$. A call to SampleRank($G, v_1, \dots, v_k, 1 - \eta$) has probability at least $1 - \eta$ of returning a correct ranking of v_1, \dots, v_k while performing $O(\frac{1}{1-\alpha} \log(\frac{4k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 p})$ queries.*

For fixed α, ϵ, k and η , SampleRank performs essentially $O(1/p)$ queries, which ranges from $O(1)$ to $O(n)$ as a function of p , and for $p \in \omega(1/n)$ is asymptotically less than the $\Omega(n)$ bound holding under strictly local exploration models (see Subsection 5.2). And note that SampleRank

relies only on *jump()* and *crawl()* queries, which are in a sense the sparsest possible global query (returning a random node of the graph), and the sparsest possible local query (returning a random child of the queried node). This provides the first theoretical justification as to why all the existing local algorithms for PageRank score or rank estimation either assume to have global access to the graph [5, 6, 11], or fail to give adequate worst-case cost guarantees [2, 7, 14, 16].

The next subsection will show that, even under any other exploration model, no algorithm can perform better than SampleRank, at least as long as the scores do not fall under $\Theta(n^{-2/3})$.

5.4 Lower Bounds for All Graph Exploration Models

This subsection presents lower bounds for Monte Carlo local algorithms that hold under every graph exploration model allowing any combination of local and global queries. These are the first lower bounds at this level of generality that we are aware of. Formally, we prove that:

Theorem 10. *Choose an integer $k \geq 2$, a damping factor $\alpha \in (0, 1)$, a separation $\epsilon > 0$, and a score function $\Theta(1/x) \leq p(x) \leq \Theta(1)$. Under any graph exploration model, for any Monte Carlo local ranking algorithm MC with confidence $\frac{1}{k!} + \delta$ there exists a family of graphs $\{G_m\}$ where G_m has size $n = \Theta(m)$ and contains k nodes v_1, \dots, v_k such that:*

1. $P(v_1), \dots, P(v_k) \in \Theta(p(n))$
2. v_1, \dots, v_k are $\approx \epsilon$ -separated, i.e. $P(v_{i+1})/P(v_i) \approx 1 + \epsilon$
3. MC performs $\Omega(\delta \min\{1/p(n), n^{2/3}\})$ queries on the instance $(G, \{v_1, \dots, v_k\})$.

By Theorem 10, our algorithm SampleRank, although based on the spare *jump()* and *crawl()* queries, is asymptotically optimal also among the class of algorithms employing far more informative ones, save possibly for PageRank scores in $o(n^{-2/3})$. And, thus, the jump and crawl model appears to be the most powerful possible for local PageRank computations, except perhaps for extremely small scores: in this latter case it remains to decide whether the discrepancy is an artifact of our analysis, or if more informative global and local queries allow to obtain better (and fully sublinear?) worst-case cost guarantees.

6 Conclusions

In the colourful universe of results and techniques for local PageRank computations, we analyse the interplay between algorithms and graph exploration models, identifying the boundaries between reasonable and prohibitive exploration costs. We show that one must combine a non-local exploration and a positive probability of error to avoid exploring essentially the whole graph in the worst case; and that if both these conditions are satisfied, a spare model providing minimum amounts of local and global access is as powerful as possible, except perhaps for extremely small PageRank scores, and allows to obtain the solution by visiting a limited portion of the graph.

For the practitioner, our results imply that it is impossible to ensure correctness when computing the PageRank score or ranking in graphs that (as the web or many social networks) are partially inaccessible or evolve rapidly in time, even if allowing for large approximation factors; and they suggest great care when designing access models (e.g. in graph manipulation libraries or social networks APIs) that should support local PageRank computations – either you include a global exploration primitive, or you cut out anyone who attempts efficient local computations.

For the theoretician, our work delineates the general framework of local PageRank computation, leaving open the only problem to determine if even the smallest PageRank scores can be computed locally for a sublinear cost, or if this apparent possibility is an artifact of our lower bounds. If the first case holds, we have shown that the only possibility to achieve a sublinear cost resides in using local and global exploration primitives more informative than the ones made available by the simple jump and crawl model.

References

- [1] A. Anagnostopoulos, R. Kumar, M. Mahdian, E. Upfal, and F. Vandin. Algorithms on evolving graphs. In *Proc. of ITCS*, pages 149–160. 2012.
- [2] R. Andersen, C. Borgs, J. Chayes, J. Hopcroft, V. Mirrokni, and S.-H. Teng. Local computation of PageRank contributions. *Internet Mathematics*, 5(1–2):23–45, 2008.
- [3] R. Andersen, F. Chung, and K. Lang. Local graph partitioning using PageRank vectors. In *Proc. of IEEE FOCS*, pages 475–486. 2006.
- [4] K. Avrachenkov, N. Litvak, D. Nemirovsky, and N. Osipova. Monte Carlo methods in PageRank computation: When one iteration is sufficient. *SIAM Journal on Numerical Analysis*, 45(2):890–904, 2007.
- [5] B. Bahmani, A. Chowdhury, and A. Goel. Fast incremental and personalized PageRank. In *Proc. of VLDB Endowment*, 4(3):173–184, 2010.
- [6] B. Bahmani, R. Kumar, M. Mahdian, and E. Upfal. Pagerank on an evolving graph. In *Proc. of ACM KDD*, pages 24–32. 2012.
- [7] Z. Bar-Yossef and L.-T. Mashiach. Local approximation of PageRank and reverse PageRank. In *Proc. of ACM CIKM*, pages 279–288. 2008.
- [8] P. Boldi, R. Posenato, M. Santini, and S. Vigna. Traps and pitfalls of topic-biased PageRank. In *Proc. of WAW*, pages 107–116. 2006.
- [9] C. Borgs, M. Brautbar, J. T. Chayes, S. Khanna, and B. Lucier. The power of local information in social networks. In *Proc. of WINE*, pages 406–419. 2012.
- [10] C. Borgs, M. Brautbar, J. T. Chayes, S. Khanna, and B. Lucier. The power of local information in social networks. *CoRR*, abs/1202.6033, 2012.
- [11] C. Borgs, M. Brautbar, J. T. Chayes, and S.-H. Teng. A sublinear time algorithm for PageRank computations. In *Proc. of WAW*, pages 41–53. 2012.
- [12] M. Brautbar and M. Kearns. Local algorithms for finding interesting individuals in large networks. In *Proc. of ICS*, pages 188–199. 2010.
- [13] M. Bressan, E. Peserico and L. Pretto. The power of local information in PageRank. In *Proc. of WWW (Companion Volume)*, pages 179–180. 2013.
- [14] M. Bressan and L. Pretto. Local computation of PageRank: the ranking side. In *Proc. of ACM CIKM*, pages 631–640. 2011.
- [15] S. Brin and L. Page. The anatomy of a large scale hypertextual Web search engine. In *Proc. of WWW*. 1998.
- [16] Y.-Y. Chen, Q. Gan, and T. Suel. Local methods for estimating PageRank values. In *Proc. of ACM CIKM*, pages 381–389. 2004.
- [17] J. Cho, H. Garcia-Molina, and L. Page. Efficient crawling through URL ordering. *Computer Networks*, 30(1–7):161–172, 1998.
- [18] X. Deng and C. H. Papadimitriou. Exploring an unknown graph (extended abstract). In *Proc. of IEEE FOCS*, pages 355–361. 1990.

- [19] D. Fogaras, B. Rácz, K. Csalogány, and T. Sarlós. Towards scaling fully personalized PageRank: Algorithms, lower bounds, and experiments. *Internet Mathematics*, 2(3):333–358, 2005.
- [20] Z. Gyöngyi, H. Garcia-Molina, and J. Pedersen. Combating web spam with TrustRank. In *Proc. of VLDB*, pages 576–587. 2004.
- [21] T. H. Haveliwala. Topic-Sensitive PageRank. In *Proc. of WWW*, pages 517–526. 2002.
- [22] B. Kalyanasundaram and K. Pruhs. Constructing competitive tours from local information. *Theoretical Computer Science*, 130(1):125–138, 1994.
- [23] A. N. Langville and C. D. Meyer. Deeper inside PageRank. *Internet Mathematics*, 1(3):335–380, 2004.
- [24] A. N. Langville and C. D. Meyer. *Google’s PageRank and Beyond: The Science of Search Engine Rankings*. Princeton University Press, 2006.
- [25] C. H. Papadimitriou and M. Yannakakis. Shortest paths without a map. In *Proc. of ICALP*, pages 610–620. 1989.
- [26] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proc. of ACM STOC*, pages 475–484. 1997.
- [27] D. Saez-Trumper, G. Comarela, V. Almeida, R. Baeza-Yates, and F. Benevenuto. Finding trendsetters in information networks. In *Proc. of ACM KDD*, pages 1014–1022. 2012.
- [28] X. Wu, V. Kumar, J. Ross Quinlan, J. Ghosh, Q. Yang, H. Motoda, G. J. McLachlan, A. Ng, B. Liu, P. S. Yu, Z.-H. Zhou, M. Steinbach, D. J. Hand, and D. Steinberg. Top 10 algorithms in data mining. *Knowledge and Information Systems*, 14:1–37, 2007.
- [29] S. Yan and D. Lee. Toward alternative measures for ranking venues: a case of database research community. In *Proc. of ACM/IEEE JCDL*, pages 235–244. 2007.

A Appendix

A.1 A note for Subsection 3.1

If a query can return an output that depends on the global arc structure of the graph, then it may serve as an oracle to solve global problems for a minimum cost – e.g., it could code the PageRank score of an input node into a properly crafted sequence of output nodes. To exclude this possibility, all the “reasonable” graph exploration models allow, often implicitly, only queries that do not disclose information about the global arc structure of the graph, if not for a small portion near the node given in input to the query itself.

A.2 Proof of Theorem 1

Proof. At any point during the exploration, *ALG* does not possess any information besides that provided by the queries’ output, and therefore besides that contained in H (which contains all the information collected along the whole execution). Therefore, *ALG* does not discriminate between different input graphs whose exploration can yield exactly H – i.e. between different graphs that are compatible with H . Hence *ALG* behaves identically on any graph compatible with H , querying the same nodes and returning the same result (in the case of Las Vegas randomized algorithms, this must happen in at least some execution). But if H is not a ranking subgraph for v_1, \dots, v_k , then there are two graphs G' and G'' , both compatible with H , where v_1, \dots, v_k have different rankings; and therefore the ranking returned by *ALG* must be wrong for at least one of the two instances $(G', \{v_1, \dots, v_k\})$ and $(G'', \{v_1, \dots, v_k\})$. Thus *ALG* does not return always a correct ranking, contradicting the hypotheses. \square

A.3 Proof of Theorem 2

Proof. Suppose by contradiction that there exists an algorithm *ALG* that obeys a model satisfying Property 1 and that solves (at least in some execution) the instance $(G, \{v_1, \dots, v_k\})$ performing $r < r^*$ queries. Consider the set formed by any node that either has been returned by a global query issued by *ALG*, or either is that on which the output of a local query issued by *ALG* depends (see Property 1); this set has then size at most r . Consider the ranking subgraph compatible with G that has this set as kernel set. Clearly, H is a supergraph of the graph witnessed by *ALG*, and its kernel nodes are known to have no neighbours outside the visit subgraph itself – thus it contains at least all the information possessed by *ALG*. All the input graphs compatible with H must then be “compatible” also with the visit performed by *ALG*, i.e., the queries’ output witnessed by *ALG* can derive from any of them with positive probability. H is not a ranking subgraph for v_1, \dots, v_k on G , since its kernel size is smaller than r^* ; therefore it is compatible with two graphs where the target nodes have different rankings. These two graphs must then be compatible also with the visit performed by *ALG*. By the same argument of the proof of Theorem 1, *ALG* does not always return a correct solution, leading to a contradiction. \square

A.4 Proof of Theorem 3

Proof. We first show that the two conditions are both necessary.

For 1), we show that if $P_H(u) < P_H(v)/(1 + \epsilon)$ then there is a graph G compatible with H where $P(v) > (1 + \epsilon)P(u)$, hence $v \succ u$ and H is not a ranking subgraph for $u \succ v$. G is built from H by adding n nodes that are pointed by every frontier node of H (if the frontier is empty, then these nodes are orphans). The scores $P(u)$ and $P(v)$ of u and v in G are the sum of two components: one accounting for paths that contain only kernel nodes, the other accounting for paths that include at least one frontier node. The first component amounts to $P_H(\cdot)|H|/|G|$, since paths containing only kernel nodes are the same in H and

G , but in G the probability of restarting at any given node is $1/|G|$ instead of $1/|H|$. The second component is the sum of contributions of at most $|H|$ different nodes, since this is the number of ancestors of u or v in H and, by construction, also in G ; and these contributions are provided via paths each containing some frontier node x , which has a fraction at most $|H|/(|H| + n) = |H|/|G|$ of children leading to the target node (the others leading to the additional n sinks). Taking into account the normalization factor $(1 - \alpha)/|G|$, the second component is bounded by $|H| \cdot (1 - \alpha)/|G| \cdot |H|/|G| < |H|^2/|G|^2$. Therefore we have:

$$\frac{P(u)}{P(v)} \leq \frac{P_H(u)|H|/|G| + |H|^2/|G|^2}{P_H(v)|H|/|G|} = \frac{P_H(u)}{P_H(v)} + \frac{|H|}{P_H(v)|G|} \quad (4)$$

but by hypothesis $P_H(u)/P_H(v) < 1 + \epsilon$, and thus choosing a sufficiently large n (and therefore $|G|$) will make $P(v) > (1 + \epsilon)P(u)$, implying that $v \succ u$ in G and that H is not a ranking subgraph for $u \succ v$.

For 2), if there exists $w \in \text{frontier}(H)$ such that $\sum_{(w,z) \in H} P_H(z, u) < \sum_{(w,z) \in H} P_H(z, v)$, then again there is a graph G compatible with H where $P(v) > (1 + \epsilon)P(u)$, hence $v \succ u$ and H is not a ranking subgraph for $u \succ v$. G is obtained from H by adding n parents to w . In G , we have $P(w, v) > P(w, u)$ since by hypothesis the children of w provide, on average, strictly more contribution to v than to u . Therefore, each of the additional parents of w also provides strictly more contribution to v than to u . Then for n sufficiently large we have again $P(v) > (1 + \epsilon)P(u)$, implying that $v \succ u$ in G and that H is not a ranking subgraph for $u \succ v$.

To prove that the two conditions are sufficient, consider any graph G compatible with H and consider again the two components of the contributions provided via paths that contain respectively only kernel nodes and at least one frontier node. As discussed above, the paths related to the first component still exist in G and provide exactly the same contributions, except for a scaling factor $|H|/|G|$; therefore, by condition 1) their overall contribution towards u is at least $1/(1 + \epsilon)$ times that towards v . The paths related to the second component must contain a frontier node w , and condition 2) ensures that, of the contributions flowing through w , the component directed to u is greater or equal than (and thus at least $1/(1 + \epsilon)$ times) that directed to v . Summing the two components proves that in G we still have $P(u) \geq P(v)/(1 + \epsilon)$, and therefore either we have a tie, which we can break as $u \succ v$, or we have $P(u) > (1 + \epsilon)P(v)$, which already gives $u \succ v$. In any case, the ranking is $u \succ v$ and H is a ranking subgraph. \square

A.5 Union of visit subgraphs

Definition 5. Let $H_1 = (V_1, A_1)$ and $H_2 = (V_2, A_2)$ be two visit subgraphs. Their union $H_1 \cup H_2$ is $H = (V, A)$ such that $V = V_1 \cup V_2$, $A = A_1 \cup A_2$, and $\text{kernel}(H) = \text{kernel}(H_1) \cup \text{kernel}(H_2)$.

Lemma 4. Let $H = H_1 \cup H_2$ be the union of two visit subgraphs H_1 and H_2 . Then

1. H is a visit subgraph
2. if H_1 and H_2 are compatible with a subgraph G , then H is also compatible with G
3. if H_1 is a ranking subgraph for two nodes u, v , then H is also a ranking subgraph for u, v

Proof. 1. If $x \in \text{frontier}(H)$, then $x \in \text{frontier}(H_1) \cup \text{frontier}(H_2)$ and $x \notin \text{kernel}(H_1) \cup \text{kernel}(H_2)$. From the first follows that x has at least one arc from/to $\text{kernel}(H_1)$ or $\text{kernel}(H_2)$ and therefore from/to $\text{kernel}(H)$. From the second follows that x has no arcs from/to $\text{frontier}(H_1)$ or $\text{frontier}(H_2)$, and since $\text{frontier}(H) \subseteq \text{frontier}(H_1) \cup \text{frontier}(H_2)$, x has no arcs from/to $\text{frontier}(H)$.

2. The union H of two subgraphs H_1 and H_2 of G is clearly a subgraph of G . Consider $x, y \in \text{kernel}(H) = \text{kernel}(H_1) \cup \text{kernel}(H_2)$; by definition, if and only if (x, y) is an arc of G then it is also an arc in H_1 and H_2 (although x or y may be frontier nodes in one of the

two) and therefore in H . When $x \in \text{kernel}(H)$ and $y \in \text{frontier}(H)$, the same reasoning proves that the nodes pointing from/to the kernel of H in G , and the relative arcs, are the same in H and in G .

3. Suppose instead that there exist two graphs G' and G'' compatible with H where u and v have different rankings. It is easy to see that G' and G'' are compatible also with H_1 , which therefore would not be a ranking subgraph for u and v , contradicting the hypothesis. \square

Lemma 5. *Let $v_1 \succ \dots \succ v_k$ be nodes of a graph G . Then the set of all ranking subgraphs for v_1, \dots, v_k compatible with G coincides with the set of all possible unions $\bigcup_{i=1}^{k-1} R_i$, where R_i is a ranking subgraph for v_i, v_{i+1} compatible with G .*

Proof. Let R be a ranking subgraph for v_1, \dots, v_k compatible with G . Then R is also a ranking subgraph for $v_i, v_{i+1} \forall i = 1, \dots, k-1$, and we can express R as $\bigcup_{i=1}^{k-1} R_i$ with $R_i = R$.

On the other side, let $S = \bigcup_{i=1}^{k-1} R_i$ where R_i is a ranking subgraph for v_i, v_{i+1} compatible with G . By Lemma 4, S is compatible with G and is still a ranking subgraph for $v_i, v_{i+1} \forall i = 1, \dots, k-1$. Therefore, it is a ranking subgraph also for v_1, \dots, v_k . \square

A.6 Proof of Theorem 4

Proof. We build a graph G_m of size $n = \Theta(m)$ satisfying the statement for every m sufficiently large. G_m is formed by k disjoint subgraphs G_m^1, \dots, G_m^k . Subgraph G_m^i (Figure 2) contains:

1. the target node v_i and its self-loop
2. $n_i = m p(m)(1 + \epsilon(m))^{i-1}$ “bulk parents” (rounded to the nearest positive integer) having v_i as their sole child; if $n_k = n_{k-1}$, then v_k has $n_k + 1$ bulk parents
3. for $i < k$, a clique of m nodes, one of which has also an arc to v_i
4. for $1 < i < k$, if $n_i = n_{i-1}$, a “fractional parent” of v_i having an arc to v_i and $2 \leq q_i \leq m-1$ arcs (we compute q_i below) to clique nodes that do not point to v_i




Figure 2: Subgraph G_m^i of graph G_m (Theorem 4) for $i < k$. The n_i “bulk parents” on the left provide the bulk of v_i ’s score, up to a resolution of $\Theta(1/n)$; if still $n_i = n_{i-1}$, then the “fractional parent” above v_i provides a score separation up to a resolution of $\Theta(1/n^2)$. The clique on the right provides a negligible contribution $O(1/n^2)$, but its $m = \Theta(n)$ nodes must still be visited.

Intuitively, v_i and its n_i bulk parents provide the bulk of v_i ’s score, guaranteeing an absolute score separation up to a “resolution” $\Omega(1/n)$. If the desired absolute separation is smaller, then $n_i = n_{i-1}$, in which case the fractional parent provides an additional contribution $\Theta(1/nq_i)$, thus

guaranteeing an absolute score separation with resolution up to $\Theta(1/n^2)$. The clique provides a negligible contribution $O(1/n^2)$ while damping the contribution of the fractional parent and guaranteeing an $\Omega(n)$ lower bound.

Disregarding the at most $2k$ target nodes and fractional parents, G_m has size:

$$n \approx \sum_{i=1}^k (n_i + m) = km + mp(m) \sum_{i=1}^k (1 + \epsilon(m))^{i-1} \quad (5)$$

since $\epsilon(m) \in O(1/p(m)^{\frac{1}{k-1}})$, this is in $O(1/p(m))$, proving that $n \in \Theta(m)$. We thus use n instead of m from now on, all the results still holding asymptotically.

The PageRank score of v_i is the sum of three components. The first, provided by v_i and its n_i bulk parents, amounts to $\Theta(\frac{1}{n}\alpha n_i) = \Theta(p(n)(1 + \epsilon(n))^{i-1})$. The second, provided by the clique for $i < k$, can be bounded observing that in an isolated clique (without incoming or outgoing arcs) each node has score exactly $1/n$, and that adding an outgoing arc does not increase this score; then expressing $P(v_i)$ as the weighted sum of the scores of its parents divided by their outdegrees shows that the whole clique contribution is roughly the score of one of its nodes divided by m , and is thus bounded by $1/nm \in O(1/n^2) = o(p(n))$. The third component is provided by the fractional parent when $n_i = n_{i-1}$ with $1 < i < k$; of this, the part flowing through the clique is damped by a factor $O(1/n^2)$ and is thus negligible, while the part flowing through the single arc to v_i amounts to $O(1/(nq_i)) = O(p(n))$. It follows that $P(v_1) \in \Theta(p(n))$ and $P(v_i) \in \Omega(p(n))$ for every $i = 1, \dots, k$.

We prove that v_1, \dots, v_k are $\Theta(\epsilon(n))$ -separated. Consider v_i and v_{i+1} . If $n_i < n_{i+1}$, then the absolute score separation between v_{i+1} and v_i is guaranteed by their bulk parents since $n_{i+1}/n_i \approx 1 + \epsilon(n)$ and thus $P(v_{i+1})/P(v_i) \approx 1 + \epsilon(n)$ – the clique contribution is negligible and the fractional parent contribution is at most roughly half that of one bulk parent. If $n_{i+1} = n_i$, the absolute score separation is between $\Omega(1/n^2)$ and $O(1/n)$. Then choose an appropriate $2 \leq q_{i+1} \leq \Theta(n)$ to approximate the separation up to a small constant factor (note that the contribution flowing to v_{i+1} through the clique is always in $O(1/n^2)$ and thus dominated by the contribution flowing through the arc pointing directly to v_{i+1}). In any case, the target nodes are $\Theta(\epsilon(n))$ -separated.

Note that v_1, \dots, v_k are the top k ranking nodes in G_m , as they receive the contribution of at least one bulk parent plus that of the fractional parent. The cliques contain the only other non-orphan nodes, which receive a contribution of at most $1/n$ from the clique itself (less than a bulk parent) and from the fractional parent at most the same contribution received by the corresponding target node.

We conclude proving the lower bound. Any ranking subgraph for v_1, \dots, v_k in G_m must contain in its kernel set all the nodes of G_m^1, \dots, G_m^{k-1} , otherwise violating the definition of ranking subgraph (if target nodes are not in the kernel set) or condition 2 of Theorem 3. It must also contain $\Omega(n_{k-1}/(1 + \epsilon(n)))$ ancestors of v_k , otherwise violating condition 1 of Theorem 3. Therefore, at most $O(n_k - n_{k-1}/(1 + \epsilon(n))) = O(n_k \epsilon(n)/(1 + \epsilon(n)))$ nodes of G_m may not be in the kernel set. But $O(n_k \epsilon(n)/(1 + \epsilon(n))) = O(np(n)\epsilon(n)(1 + \epsilon(n))^{k-2}) = nO(p(n)\epsilon(n)(1 + \epsilon(n))^{k-2})$, and therefore the kernel set has size at least $n(1 - O(p(n)\epsilon(n)(1 + \epsilon(n))^{k-2}))$. By Theorem 2, any execution of any deterministic or randomized Las Vegas local ranking algorithm under any graph exploration model needs $n(1 - O(p(n)\epsilon(n)(1 + \epsilon(n))^{k-2}))$ queries to solve the instance $(G_m, \{v_1, \dots, v_k\})$. \square

A.7 Proof of Theorem 5

Proof. CLIQUE asks if a directed graph G_0 contains a clique of size m (the undirected version is straightforwardly reduced to this one). Given a generic CLIQUE instance (G_0, m) , we build an instance for RANKGRAPH that admits a ranking subgraph of size $2(m + 1) + q$, with q depending on ϵ , if and only if G_0 contains a clique of size m (see Figure 3).

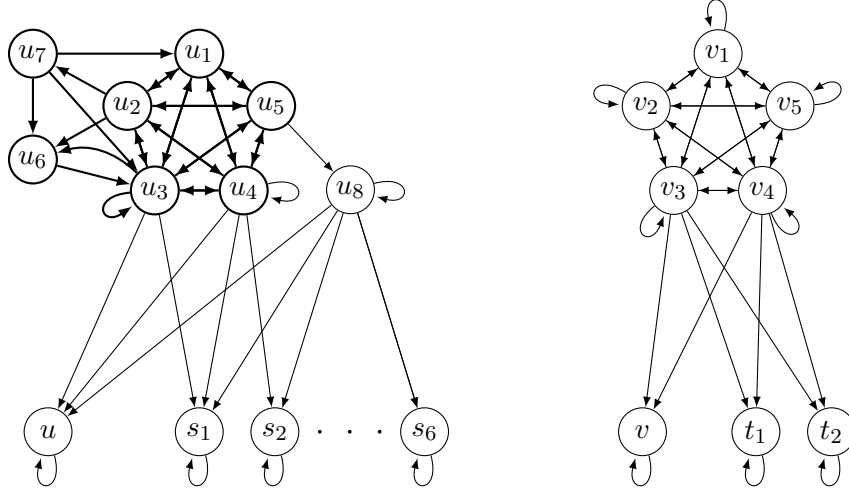


Figure 3: The graph G built from a CLIQUE instance $(G_0, 5)$. The graph G_0 contains the nodes and arcs in bold; all the others are added by the reduction (for simplicity, of the arcs added to G_0 we show only those added to u_3, u_4 and u_8 , and of the arcs going to v, t_1 and t_2 we show only those from v_3 and v_4). Here u_2 has outdegree 6 in G_0 , and thus $d = 8$. Since $u \succ v$, there exists a ranking subgraph of size 12 for u, v in G if and only if G_0 contains a clique of size 5 (for simplicity, we assumed ϵ sufficiently small that $q = 0$, i.e. v does not have additional parents besides the clique nodes).

We build the ancestor graph of u starting from G_0 . First add a self-loop to u . Then, to each node of G_0 we add an arc towards u , and a self-loop if not already present (this “normalizes” the topology with respect to already present self-loops). Let d be the maximum outdegree of a node after these modifications; then bring to d the outdegree of each node by adding a proper number of arcs towards a set of (at most $d - 2$) additional new nodes that act as sinks. The ancestor graph of v contains v with a self-loop, and a clique of m nodes each of which has a self-loop, an arc to v , and $d - m - 1$ arcs towards a set of $d - m - 1$ additional nodes that act as sinks (if $d < m + 1$ then G_0 has maximum outdegree $< m$, in which case we immediately return NO to CLIQUE without solving RANKGRAPH). In addition, v has a set of q parents each having outdegree $\lceil d^2/\alpha(1 + \epsilon) \rceil$, pointing to v and to $\lceil d^2/\alpha(1 + \epsilon) \rceil - 1$ sink nodes. We choose q such that, for $|H| = 2(m + 1) + q$, the kernel score component $P_H^Q(v)$ provided by the q parents satisfies $\epsilon P_H^C(v) - \alpha^2(1 + \epsilon)/|H|d^2 \leq P_H^Q(v) \leq \epsilon P_H^C(v)$ where $P_H^C(v)$ is the component provided by v and its m clique ancestors. A q satisfying this constraint always exists (and it equals 0 for ϵ sufficiently small), since a single parent gives contribution at most $\alpha^2(1 + \epsilon)/|H|d^2$, ensuring sufficient “resolution”. Since clearly $P_H(v) = P_H^C(v) + P_H^Q(v)$, we have $(1 + \epsilon)(P_H^C(v) - \alpha^2/|H|d^2) \leq P_H(v) \leq (1 + \epsilon)P_H^C(v)$.

Denote by G the resulting graph, and consider the input instance $(G, u, v, 2(m + 1) + q)$ for RANKGRAPH. The intuition is that a ranking subgraph of kernel size $2(m + 1) + q$ exists if and only if there are m ancestors of u that induce a subgraph “strong enough” to guarantee $P_H(u) \geq P_H(v)/(1 + \epsilon)$ – and this happens exactly if they form a clique. Indeed we prove:

1. If G_0 contains a clique of size m , then G has a compatible ranking subgraph of size $2(m + 1) + q$ for $u \succ v$. Indeed, consider the visit subgraph having in the kernel set the m clique nodes, along with u, v , and all the $m + q$ ancestors of v . This visit subgraph has kernel size $2(m + 1) + q$ and also satisfies property 2 of Theorem 3. Observe that $P_H(u) = P_H^C(v)$ since the subgraphs induced by u with its m ancestors and by v with its m clique ancestors are identical – cliques of m nodes with outdegree d , each having in addition a self-loop and an arc to the target node. But from above we have $P_H^C(v) \geq P_H(v)/(1 + \epsilon)$, and therefore $P_H(u) \geq P_H(v)/(1 + \epsilon)$; thus H satisfies also property 1 of Theorem 3 and

is a ranking subgraph of kernel size $2(m+1) + q$ for $u \succ v$ compatible with G .

2. If G_0 does not contain a clique of size m , then either $P(u) < P(v)/(1+\epsilon)$ in G (which we can check by computing their scores, and in which case we immediately return NO to CLIQUE without solving RANKGRAPH) or there is no ranking subgraph of kernel size $2(m+1) + q$ for $u \succ v$ in G , as we prove. A ranking subgraph H for $u \succ v$ must include u and v by definition, and the $m+q$ ancestors of v by property 2 of Theorem 3; thus, to have kernel size $2(m+1) + q$, it must contain at most m ancestors of u . If these m ancestors do not induce a clique, at least one arc is missing between two of them (it instead points to a frontier node, and any path using it gives null kernel contribution). If these nodes formed a clique, this arc would provide a kernel contribution via paths starting with it and leading to u via another arc (followed by zero or more self-loops) amounting to $\alpha^2/|H|d^2$; and additional contribution via paths that use it as intermediate arc. The kernel score of u in this “clique” scenario is $P_H^C(u) = P_H^C(v)$, and thus we have $P_H(u) < P_H^C(v) - \alpha^2/|H|d^2$. With the inequalities above this implies $P_H(u) < P(v)/(1+\epsilon)$, inducing $v \succ u$ and proving that H is not a ranking subgraph for $u \succ v$.

Therefore we have a (clearly polynomial) reduction from CLIQUE to RANKGRAPH, which concludes the proof. \square

A.8 Proof of Theorem 6

Proof. Given a graph G_0 and an integer $m > 0$, DOMINATING SET asks if there is a subset D of m nodes in G_0 such that any other node in G_0 points to a node in D (the undirected version can be straightforwardly reduced to this one).

Starting from a generic directed DOMINATING SET instance (G_0, m) , we build a graph G (Figure 4) such that a compatible ranking subgraph of size $3 + m$ exists for $u \succ v$ if and only if G_0 contains a dominating set of size m . We assume that G_0 has no dangling nodes; otherwise, the dominating set is formed by all the dangling nodes plus the dominating set on the remaining nodes. Denote the nodes of G_0 by v_1, \dots, v_n . The graph G contains u and v




Figure 4: The graph G built from a DOMINATING SET instance $(G_0, 2)$. The graph G_0 contains 5 nodes v_1, \dots, v_5 and the arc (u'_i, u_j) exists in G if and only if (v_i, v_j) exists in G_0 . Graph G admits a compatible ranking subgraph with kernel size 5 if and only if G_0 admits a dominating set of size 2. In this example, such a ranking subgraph has u, u_1, u_3, v, w as its kernel set, and thus a dominating set of size 2 for G_0 is $\{v_1, v_3\}$.

with their self-loops; n parents u_1, \dots, u_n of u ; one parent w of v ; and n nodes u'_1, \dots, u'_n where u'_i has one arc towards w and, for every $j = 1, \dots, n$, an arc towards u_j if and only if the arc (v_i, v_j) exists in G_0 . The intuition is that a ranking subgraph of kernel size $3 + m$ exists if and

only if there are m parents of u that, if in the kernel set, can “cover” the frontier nodes among u'_1, \dots, u'_n ensuring that none of them gives more contribution to v than to u . Indeed we prove:

1. If G_0 contains a dominating set of size m , then there exists a ranking subgraph of size $3 + m$ for $u \succ v$ in G . Let D be the dominating set of size m in G_0 . Consider the visit subgraph H having in its kernel set u, v, w and, for $i = 1, \dots, n$, node u_i if and only if v_i is in the dominating set, thus having all the other nodes of G as frontier nodes. Since $m > 0$, then u has at least one parent in the kernel set and $P_H(u) \geq P_H(v)$; therefore H satisfies property 1 of Theorem 3 for any $\epsilon \geq 0$. Since each node u'_i points to at least a parent of u and to exactly one parent of v , then H also satisfies property 2 of Theorem 3, proving that it is a ranking subgraph of kernel size $3 + m$ for $u \succ v$ compatible with G .
2. If there exists a ranking subgraph of size $3 + m$ for $u \succ v$ in G , then we can build a dominating set of size m for G_0 as follows. The kernel set of any ranking subgraph for $u \succ v$ in G contains v and w by property 2 of Theorem 3, and contains u by definition. Therefore, if the kernel set has size $3 + m$, it must include exactly m nodes among $u_1, \dots, u_n, u'_1, \dots, u'_n$. Again by property 2 of Theorem 3, either the generic node u'_i is in the kernel set, or one of its children u_j is in the kernel set, otherwise $P_H(u'_i, u) = 0 < P_H(u'_i, v)$. For each u'_i in the kernel set, we replace it with one of its children (even if already in the kernel set); this does not increase the kernel size and maintains property 2 of Theorem 3. We end up with a visit subgraph whose kernel set contains $m' \leq m$ nodes among u_1, \dots, u_n that “cover” (i.e. have an arc from) each of the frontier nodes u'_1, \dots, u'_n . A dominating set of size m for G_0 contains v_i if and only if u_i is in the kernel set, and $m - m'$ other nodes in G_0 .

Therefore we have a (clearly polynomial) reduction from DOMINATING SET to RANKGRAPH, which concludes the proof. \square

A.9 Proof of Lemma 1

Proof. Given an instance (G_0, m) for DOMINATING SET, the reduction builds an instance $(G, u, v, 3 + m)$ for RANKGRAPH where $|G| = 2|G_0| + 3$. Furthermore, a solution of size $3 + m$ for RANKGRAPH exists if and only if there is a solution of size m for DOMINATING SET. This implies that approximating RANKGRAPH within a small (surely up to polynomial in $|G|$) factor would allow approximating DOMINATING SET within the same factor. Approximating DOMINATING SET within $o(\log(|G_0|))$ is NP-hard [26], thus approximating RANKGRAPH within $o(\log(|G|))$ is also NP-hard. \square

A.10 Proof of Theorem 7

Proof. Let $m = \lceil 1/\eta \rceil$. The graph G is formed by u and v , $m + 1$ orphan nodes that point only to u , and m orphan nodes that point only to v . Therefore, $P(u) > P(v)$. The algorithm, EludeRS, is given in the next page.

It is immediate to verify that EludeRS on $(G, \{u, v\})$ correctly returns $u \succ v$ without querying all the parents of v . By Theorem 3, all the parents of v must belong to the kernel set of any ranking subgraph for u, v in G . Therefore, EludeRS never visits such a ranking subgraph. It is also immediate to verify that, on any instance containing a graph different from G , EludeRS visits all the ancestors of u and v and returns the correct ranking; unless u and v have zero children and respectively $m + 1$ and m parents, all these parents are orphans and have one child with the exception of one of v 's parents, and this parent is exactly that not queried by EludeRS. This happens with probability at most $1/m \leq \eta$, and thus EludeRS has confidence $1 - \eta$.

Note that the score separation is $\epsilon = \alpha/(1 + \alpha m)$ in the simple instance above, but any value of ϵ can be approximated by adding parents to u or (with additional children) to v , and tuning EludeRS accordingly. \square

Algorithm 3 *EludeRS*($G, \{u, v\}$)

```
Call links() on  $u$  and  $v$ 
if ( $u$  has  $m + 1$  parents and 0 children) and ( $v$  has  $m$  parents and 0 children) then
    Call links() on all the parents of  $u$ 
    Call links() on all the parents of  $v$  except one chosen uniformly at random
    if (all the queried parents are orphans and have one child) then
        return  $u \succ v$ 
    else
        Visit the entire ancestor graph of  $u$  and  $v$  and return the correct ranking
    end if
else
    Visit the entire ancestor graph of  $u$  and  $v$  and return the correct ranking
end if
```

A.11 Proof of Theorem 8

Proof. We build a graph G_m of size $n = \Theta(m)$ satisfying the statement for every m sufficiently large. G_m consists of k disjoint subgraphs G_m^1, \dots, G_m^k . Subgraph G_m^i (Figure 5) contains the target node v_i and its self-loop, a chain of $\log_\alpha p(m)$ nodes ending with v_i , and m nodes having the head of the chain as their sole child; one of them, called “strong ancestor”, has $((1 + \epsilon)^{i-1} - 1)m/\alpha$ exclusive parents. All these quantities are meant as rounded to the nearest integer – we assume m to be sufficiently large to avoid dealing with rounding operators and still guarantee the asymptotic results of the statement. Intuitively, MC can only explore the graph “from below” to find at least one strong ancestor, which happens with probability proportional to the fraction of nodes visited, otherwise the returned ranking is correct with probability at most $1/k!$.




Figure 5: Subgraph G_m^i of graph G_m (Theorem 8). The bulk of v_i ’s PageRank score is provided by its m ancestors via a chain of $\log_\alpha p(m)$ nodes (damping their contribution by a factor $p(m)$), while its relative ranking is precisely determined by the number of parents of its “strong ancestor” (filled in grey). Any (Monte Carlo) algorithm without global queries can at best crawl backward from v_i and sample the m ancestors among which lies the strong ancestor.

G_m has size:

$$n = \sum_{i=1}^k \left(\log_{\alpha} p(m) + m + m \frac{(1+\epsilon)^{i-1} - 1}{\alpha} \right) \quad (6)$$

which is in $\Theta(m)$ since $\log_{\alpha} p(m) \in O(\log_{\alpha}(1/m)) = O(m)$ and $\alpha, \epsilon, k \in \Theta(1)$. The PageRank score of v_i is

$$P(v_i) = \frac{1}{n} \left(\sum_{j=0}^{\log_{\alpha} p(m)-1} \alpha^j + \alpha^{\log_{\alpha} p(m)} \left(m + \alpha \frac{(1+\epsilon)^{i-1} - 1}{\alpha} m \right) \right) \approx p(m)(1+\epsilon)^{i-1} \quad (7)$$

which is in $\Theta(p(n))$ since $\alpha, \epsilon, k \in \Theta(1)$ and $n = \Theta(m)$. The scores satisfy $P(v_{i+1}) \approx (1+\epsilon)P(v_i)$, making the target nodes $\approx \epsilon$ -separated (a strict ϵ -separation is obtained giving $((1+\epsilon)^{i-1} - 1)m/\alpha$ parents to v_i 's strong ancestor for some constant c , but we avoid it to keep the proof simple).

We prove the lower bound, showing that it holds even if the algorithm is aware of the structure of the graph, except for the identities and the number of parents of the strong ancestors. Since MC cannot perform global queries (which may lead to the parents of the strong ancestors or the strong ancestors themselves), it must explore the graph via local queries starting from the target nodes, “crawling backward” along the chains, and directly querying the $\Theta(m)$ ancestors, essentially sampling them (with e.g. *links()*) until finding the strong ancestor. Note that, for a non-uniform sampling, there exists a worst-case instance (having the least-queried ancestor as the strong ancestor) requiring a strictly higher number of queries than for a uniform sampling – therefore we assume that MC performs a uniform sampling. Suppose MC performs q queries overall. The probability γ that MC finds at least one strong ancestor, in which case it returns the correct ranking with probability at most 1, satisfies then $\gamma \leq q/m$. With probability $1 - \gamma$, MC does not visit any strong ancestor and must return an arbitrary ranking – the strategy minimizing the worst-case error probability consists again in returning a ranking uniformly at random among the $1/k!$ possible rankings. The overall probability of returning the correct ranking is thus upper bounded by

$$\gamma 1 + (1 - \gamma) \frac{1}{k!} \leq \frac{1}{k!} + \gamma \quad (8)$$

Since MC has confidence $1/k! + \delta$, we have $\delta \leq \gamma$, and since $\gamma \leq q/m$ we have $q \geq \delta m \in \Omega(\delta n)$. \square

A.12 Proof of Lemma 2

Proof. Note that SampleNode performs a random walk on the graph G as if it was preprocessed by adding, to each dangling node, n outgoing links to itself and to every node (this is emulated by the call to *jump()* made when *crawl()* returns an empty subgraph). The probability that SampleNode follows exactly τ arcs is $(1 - \alpha)\alpha^{\tau}$. Conditioned on this event, the probability that SampleNode returns v equals $\frac{1}{n} \sum_{z \in G} \inf_{\tau}(z, v)$ (recall from Section 2 that $\inf_{\tau}(z, v)$ denotes the probability that, starting from node z and following at each step a random outgoing link of the current node, after exactly τ steps one is at v). Then the probability that SampleNode returns v equals $\sum_{\tau=0}^{\infty} (1 - \alpha)\alpha^{\tau} \frac{1}{n} \sum_{z \in G} \inf_{\tau}(z, v)$, which coincides with $P(v)$ by Equation (1). \square

A.13 Proof of Lemma 3

Proof. Each time SampleNode goes through the first instruction of the loop, it terminates immediately with an independent probability $(1 - \alpha)$; and, by the j -th time, it has issued at most $1 + 2(j - 1) \leq 2j$ queries. Thus, SampleNode terminates after issuing at most $\frac{2}{1-\alpha}$

queries in expectation. If SampleNode issues more than $\frac{2(1+\Delta)m}{1-\alpha}$ queries, then it goes through the first instruction in the loop more than $\frac{(1+\Delta)m}{1-\alpha}$ times; the probability of this happening over m calls is the probability that that instruction makes SampleNode terminate less than $m = (1 - \frac{\Delta}{1+\Delta}) \cdot (1 + \Delta)m$ out of $\frac{(1+\Delta)m}{1-\alpha}$ times. By a simple Chernoff bound, this probability is less than $e^{-\frac{(1+\Delta)m(\frac{\Delta}{1+\Delta})^2}{2}} = e^{-\frac{m}{2} \cdot \frac{\Delta^2}{1+\Delta}}$. \square

A.14 Proof of Theorem 9

Proof. By construction, the ranking returned by SampleRank is incorrect only if some scores fall outside of their $1 - \eta/k$ confidence intervals, which by a union bound happens with probability at most $k \cdot \eta/k = \eta$. Therefore, the ranking is correct with probability at least $1 - \eta$. The rest of the proof shows that SampleRank performs more than $O(\frac{1}{1-\alpha} \log(\frac{4k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 p})$ queries with probability at most η . Precisely, we show that with probability at least $1 - \eta/2$ it performs at most $O(\log(\frac{4k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 p})$ calls to SampleNode, and that with probability at least $1 - \eta/2$ these require $O(\frac{1}{1-\alpha})$ as many queries – yielding the desired bound with probability at least $(1 - \eta/2)^2 \geq 1 - \eta$.

Consider two target nodes u, v and suppose that $P(v) = (1 + \epsilon/2)P(u)$. The intuition is that, after a sufficient number of samples, with probability at least $1 - \eta/2k$ the confidence intervals obtained from the estimates $\hat{P}_m(u)$ and $\hat{P}_m(v)$ either are disjoint (which holds also for all $P(v) \geq (1 + \epsilon/2)P(u)$) or one is within $1 + \epsilon$ of the other (which holds also for all $P(v) \leq (1 + \epsilon/2)P(u)$). We first bound the probability that the intervals become disjoint. The probability that, after m samples, a node with score at least $P(u)(1 + \epsilon/4)$ receives a score estimate smaller or equal to $P(u)(1 + \epsilon/8) = P(u)(1 + \epsilon/4)(1 - \frac{\epsilon/8}{1+\epsilon/4})$ is, by a Chernoff bound, less than $e^{-mP(u)(1+\epsilon/4)\frac{(\epsilon/8)^2}{2(1+\epsilon/4)^2}} = e^{-\Theta(mP(u)\frac{\epsilon^2}{1+\epsilon})}$. For a proper $m = O(\log(\frac{2k}{\eta}) \frac{1+\epsilon}{\epsilon^2 P(u)})$, this probability is upper bounded by $\eta/2k$, and thus if $\hat{P}_m(u) \leq P(u)(1 + \epsilon/8)$, the confidence interval for $P(v)$ cannot contain the value $P(u)(1 + \epsilon/4)$. By a Chernoff bound, $\hat{P}_m(u) > P(u)(1 + \epsilon/8)$ with probability less than $e^{-mP(u)\frac{(\epsilon/8)^2}{2+\epsilon/8}} = e^{-\Theta(mP(u)\frac{\epsilon^2}{1+\epsilon})}$, which again for a proper $m = O(\log(\frac{2k}{\eta}) \frac{1+\epsilon}{\epsilon^2 P(u)})$ is upper bounded by $\eta/2k$. Therefore, after $O(\log(\frac{2k}{\eta}) \frac{1+\epsilon}{\epsilon^2 P(u)})$ samples, with probability at least $1 - \eta/2k$ the confidence interval for $P(u)$ does not contain $P(u)(1 + \epsilon/4)$. A similar calculation shows that, after $O(\log(\frac{2k}{\eta}) \frac{1+\epsilon}{\epsilon^2 P(u)})$ samples, with probability $1 - \eta/2k$ also the confidence interval for $P(v)$ does not contain $P(u)(1 + \epsilon/4)$, and thus $C_m(u) \cap C_m(v) = \emptyset$. A similar argument shows that, after $O(\log(\frac{2k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 P(u)})$ samples, with probability at least $1 - \eta/2k$ the confidence interval of $P(u)$ does not include values below $P(u)(1 - \epsilon/4(1 + \epsilon))$ and that of $P(v)$ does not include values above $P(v)(1 + \epsilon/(8 + 4\epsilon))$. The ratio between these two extremes is exactly $1 + \epsilon$, and therefore with probability at least $1 - \eta/2k$ we have $P_m^U(v)/P_m^L(u) \leq 1 + \epsilon$. Hence, after $O(\log(\frac{2k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 P(u)})$ samples we have that, ideally ordering the nodes by their true scores, each pair of consecutive nodes has intervals that overlap or are not within $1 + \epsilon$ one of the other with probability at most $\eta/2k$. By a union bound, then, with probability $1 - \eta/2$ SampleRank terminates after $O(\log(\frac{2k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 P(u)})$ iterations (and the probability that it exceeds the expected number of iterations by a factor $c \geq 1$ decreases as e^{-c}).

By Lemma 3, the probability that $O(\log(\frac{2k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 P(u)})$ calls to SampleNode perform more than $\Theta(\frac{1+\Delta}{1-\alpha} \log(\frac{4k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 p})$ queries is bounded by $e^{-\Theta(\log(\frac{4k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 p} \frac{\Delta^2}{1+\Delta})} = O((\eta/4k) \frac{(1+\epsilon)^2 \Delta^2}{\epsilon^2 p(1+\Delta)})$, which is strictly less than $\eta/2$ for appropriate constants (and e.g. $\Delta = 2$). Therefore SampleRank performs $O(\frac{1}{1-\alpha} \log(\frac{4k}{\eta}) \frac{(1+\epsilon)^2}{\epsilon^2 p})$ with probability $(1 - \eta/2)^2 > 1 - \eta$. This concludes the proof. \square

A.15 Proof of Theorem 10

Proof. We build a graph G_m of size $n = \Theta(m)$ satisfying the statement for every m sufficiently large, first for the case $p(x) \geq \Theta(x^{-2/3})$, then extending the result to $p(x) \geq \Theta(1/x)$. G_m consists of k disjoint subgraphs G_m^1, \dots, G_m^k . Subgraph G_m^i (Figure 6) contains the target node v_i with a self-loop, and $m\sqrt{p(m)}/\alpha$ parents of v_i ; one of these, called “strong ancestor”, has v_i as its sole child and has $mp(m)((1+\epsilon)^{i-1}-1)/\alpha^2$ orphan parents, and all the others have outdegree $1/\sqrt{p(m)}$, pointing to v_i and to $1/\sqrt{p(m)}-1$ exclusive sink nodes (to keep Figure 6 legible, we have explicitly drawn these sinks only for one parent). All these quantities are meant as rounded to the nearest integer – we assume m to be sufficiently large to avoid dealing with rounding operators and still guarantee the asymptotic results of the statement.




Figure 6: Subgraph G_m^i of graph G_m (Theorem 10). The bulk of v_i 's PageRank score is provided by its $\approx m\sqrt{p(m)}$ parents, while its relative ranking is precisely determined by the number of parents of its “strong ancestor” (filled in grey). Any Monte Carlo algorithm can find a strong ancestor using either global queries to directly discover it or one of its parents (which are a fraction $O(p(n))$ of all nodes), or local queries to sample v_i 's parents (which are $\Omega(n^{2/3})$).

G_m has size:

$$\begin{aligned} n &\approx \sum_{i=1}^k \left(\frac{m\sqrt{p(m)}}{\alpha} \frac{1}{\sqrt{p(m)}} + \frac{mp(m)((1+\epsilon)^{i-1}-1)}{\alpha^2} \right) \\ &= m \left(\frac{k}{\alpha} + p(m) \sum_{i=1}^k \frac{((1+\epsilon)^{i-1}-1)}{\alpha^2} \right) \end{aligned} \quad (9)$$

which is in $\Theta(m)$ since $\epsilon, \alpha, k \in \Theta(1)$ and $p(m) \in O(1)$.

The PageRank score of v_i is:

$$\begin{aligned} P(v_i) &= \frac{1}{n} \left(1 + \alpha \left(\frac{m\sqrt{p(m)}}{\alpha} - 1 \right) \frac{1}{1/\sqrt{p(m)}} + \alpha + \alpha^2 \frac{mp(m)((1+\epsilon)^{i-1}-1)}{\alpha^2} \right) \\ &\approx p(m)(1+\epsilon)^{i-1} \end{aligned} \quad (10)$$

which is in $\Theta(p(n))$ since $\epsilon \in O(1)$ and $n = \Theta(m)$. The scores satisfy $P(v_{i+1}) \approx (1+\epsilon)P(v_i)$, making the target nodes $\approx \epsilon$ -separated (a strict ϵ -separation is obtained giving $mp(m)((1+\epsilon)^{i-1}-1)/\alpha^2$ parents to v_i 's strong ancestor for some constant c , but we avoid it to keep the proof simple).

We prove the lower bound, showing that it holds even if the algorithm is aware of the structure of the graph, except for the identities and the number of parents of the strong ancestors. To prove the bound under *any* graph exploration model, we adopt the “most powerful” local exploration model possible, which allows *jump()* and *links()* queries as we prove immediately. To find at least one strong ancestor, MC can:

1. Explore the graph using global queries (recall that global queries do not take any input and return in output a single node depending on the set of nodes of the graph, but not on its arcs). Since the output of a global query does not depend on what the algorithm has already visited, we imagine that all the global queries have been performed before all the local queries, maximizing the information obtained by MC. Observe now that $jump()$ is the “less adversarial” type of global query possible. Indeed, for any global query returning nodes not uniformly at random, we can reduce the amount of information obtained by MC by permuting the parents of a target node so that the strong ancestor and its parents are the nodes returned less likely. Even if MC is aware of this and tries to use global queries to “exclude” parents that are less likely to be strong ancestors, it needs $\Omega(n)$ queries to restrict to $O(n)$ candidates. Thus, we assume that $jump()$ is allowed. In this case, the probability that any $jump()$ query returns a strong ancestor or one of its parents (in which case the algorithm can identify the strong ancestor with only an additional $links()$ query) is in $O((np(n)((1+\epsilon)^{i-1} - 1)/\alpha^2)/n) = O(p(n))$, and for q queries it is in $O(qp(n))$.
2. Explore the graph using local queries. Note that a randomized algorithm using $links()$ can emulate any other type of local query; thus, we assume that $links()$ is allowed. MC can then start from the target nodes and sample their parents – the only ancestors that it knows besides those returned by $jump()$ – until finding the strong ancestor. For a non-uniform sampling, there exists a worst-case instance (having the least-queried ancestor as the strong ancestor) requiring a strictly higher number of queries than for a uniform sampling – therefore we assume that MC performs a uniform sampling. The probability that querying q parents yields a strong ancestor is in $O(q/n\sqrt{p(n)})$ which, since $p(n) \in \Omega(1/n^{2/3})$, is in $O(qn^{-\frac{2}{3}})$.

Therefore, if MC performs $O(q)$ queries, the probability γ of finding at least one strong ancestor, in which case it returns the correct ranking with probability at most 1, is $O(q \max\{p(n), n^{-2/3}\})$; and hence $q \in \Omega(\gamma / \max\{p(n), n^{-2/3}\})$. With probability $1 - \gamma$, MC does not find any strong ancestor, and then outputs an arbitrary ranking – again, the strategy minimizing the worst-case error probability consists in returning a ranking uniformly at random among the $k!$ possible rankings, which therefore is correct with probability at most $1/k!$. The overall probability of returning a correct ranking is thus upper bounded by

$$\gamma + (1 - \gamma) \frac{1}{k!} \leq \frac{1}{k!} + \gamma \quad (11)$$

Since MC has confidence $1/k! + \delta$, we have $\delta \leq \gamma$, and therefore $q \in \Omega(\delta / \max\{p(n), n^{-2/3}\})$.

To prove the lower bound also for $p(x) \in o(x^{-2/3})$, consider the graph built for $p(x) = x^{-2/3}$, and add a chain of $\log_\alpha(p(m)/m^{-2/3})$ nodes between the parents of v_i and v_i itself. This damps the contribution of the parents by a factor $\alpha^{\log_\alpha(p(m)/m^{-2/3})} = p(m)/m^{-2/3}$, guaranteeing $P(v_i) \in \Theta(p(n))$ while preserving the relative score separation of the target nodes. Clearly, MC solves correctly this instance if and only if it solves correctly the instance built for $p(x) = x^{-2/3}$, and the same bound holds. \square